arXiv:1407.3377v1 [cs.DS] 12 Jul 2014

Strengthened Lazy Heaps: Surpassing the Lower Bounds for Binary Heaps

Stefan Edelkamp
Faculty 3—Mathematics and Computer Science
University of Bremen
PO Box 330 440, 28334 Bremen, Germany

Amr Elmasry
Department of Computer Engineering and Systems
Alexandria University
Alexandria 21544, Egypt

Jyrki Katajainen
Department of Computer Science
University of Copenhagen
Universitetsparken 5, 2100 Copenhagen East, Denmark

July 15, 2014

Abstract

Let n denote the number of elements currently in a data structure. An *in-place heap* is stored in the first n locations of an array, uses $O(1)$ extra space, and supports the operations: *minimum*, *insert*, and *extract-min*. We introduce an in-place heap, for which *minimum* and *insert* take $O(1)$ worst-case time, and *extract-min* takes $O(\lg n)$ worst-case time and involves at most $\lg n + O(1)$ element comparisons. The achieved bounds are optimal to within additive constant terms for the number of element comparisons. In particular, these bounds for both *insert* and *extract-min*—and the time bound for *insert*—surpass the corresponding lower bounds known for binary heaps, though our data structure is similar. In a binary heap, when viewed as a nearly complete binary tree, every node other than the root obeys the heap property, i.e. the element at a node is not smaller than that at its parent. To surpass the lower bound for *extract-min*, we reinforce a stronger property at the bottom levels of the heap that the element at any right child is not smaller than that at its left sibling. To surpass the lower bound for *insert*, we buffer insertions and allow $O(\lg^2 n)$ nodes to violate heap order in relation to their parents.

1 Introduction

A *binary heap*, invented by Williams [19], is an in-place data structure that 1) implements a priority queue (i.e. supports the operations *minimum*, *construct*, *insert*, and *extract-min*); 2) requires $O(1)$ words of extra space in addition to an array storing the current collection of elements; and 3) is viewed as a nearly complete binary tree where, for every node other than the root, the element at that node is not smaller than the element at its parent (*min-heap property*). Letting n denote the number of elements in the data structure, a binary heap supports *minimum* in $O(1)$ worst-case time, and *insert* and *extract-min* in $O(\lg n)$ worst-case time. For Williams' original proposal [19], the number of element comparisons performed by *insert* is at most $\lg n + O(1)$ and that by *extract-min* is at most $2 \lg n + O(1)$. Immediately after the appearance of Williams' paper, Floyd showed [12] how to support *construct*, which builds a heap for n elements, in $O(n)$ worst-case time with at most $2n$ element comparisons.

Since a binary heap does not support all the operations optimally, many attempts have been made to develop priority queues supporting the same set (or even a larger set) of operations that improve the worst-case running time of the operations as well as the number of element comparisons performed by them [1, 3, 5, 6, 8, 10, 13, 17]. In Table 1 we summarize the fascinating history of the problem, considering the space and comparison complexities.

When talking about optimality, we have to separate three different concepts. Assume that, for a problem of size n, the bound achieved is $\mathrm{A}(n)$ and the best possible bound is $\mathrm{OPT}(n)$.

Asymptotic optimality: $\mathrm{A}(n) = O(\mathrm{OPT}(n))$.

Constant-factor optimality: $\mathrm{A}(n) = \mathrm{OPT}(n) + o(\mathrm{OPT}(n))$.

Up-to-additive-constant optimality: $\mathrm{A}(n) = \mathrm{OPT}(n) + O(1)$.

As to the amount of space used and the number of element comparisons performed, we aim at up-to-additive-constant optimality. From the information-theoretic lower bound for sorting [15, Section 5.3.1], it follows that, in the worst case, either *insert* or *extract-min* must perform at least $\lg n - O(1)$ element comparisons. As to the running times, we aim at asymptotic optimality. Our last natural goal is to support *insert* in $O(1)$ worst-case time, because then *construct* can be trivially realized in linear time by repeated insertions.

The binomial queue [17] was the first priority queue supporting *insert* in $O(1)$ worst-case time. (This was mentioned as a short note at the end of Brown's paper [1].) However, the binomial queue is a pointer-based data structure requiring $O(n)$ pointers in addition to the elements. For binary heaps, Gonnet and Munro showed [13] how to perform *insert* using at most $\lg \lg n + O(1)$ element comparisons and *extract-min* using at most $\lg n + \lg^* n + O(1)$ element comparisons. Carlsson et al. showed [5] how to achieve $O(1)$ worst-case time per *insert* by an in-place data structure that utilizes a queue of pennants. (A pennant is a binary heap with an extra root that has one child.) For this data structure,

Table 1: The worst-case comparison complexity of some priority queues. Here n denotes the number of elements stored, w is the size of machine words in bits, and the amount of extra space is measured in words. All these data structures support *construct* in $O(n)$ worst-case time and *minimum* in $O(1)$ worst-case time.

Data structure	**Space**	*insert*	*extract-min*
binary heaps [19]	$O(1)$	$\lg n + O(1)$	$2 \lg n + O(1)$
binomial queues [1, 17]	$O(n)$	$O(1)$	$2 \lg n + O(1)$
heaps on heaps [13]	$O(1)$	$\lg \lg n + O(1)$	$\lg n + \log^* n + O(1)$
queue of pennants [5]	$O(1)$	$O(1)$	$3 \lg n + \log^* n + O(1)$
multipartite priority queues [10]	$O(n)$	$O(1)$	$\lg n + O(1)$
engineered weak heaps [8]	$n/w + O(1)$	$O(1)$	$\lg n + O(1)$
strengthened lazy heaps [this work]	$O(1)$	$O(1)$	$\lg n + O(1)$

the number of element comparisons performed per *extract-min* is bounded by $3 \lg n + \log^* n + O(1)$. The multipartite priority queue [10] was the first priority queue achieving the asymptotically-optimal time and additive-term-optimal comparison bounds: $O(1)$ worst-case time for *minimum* and *insert*, and $O(\lg n)$ worst-case time with at most $\lg n + O(1)$ element comparisons for *extract-min*. Unfortunately, the structure is involved and its representation requires $O(n)$ pointers. Another solution by us [8] is based on weak heaps [7]. To implement *insert* in $O(1)$ worst-case time, we use a bulk-insertion strategy—employing two buffers and incrementally merging one with the weak heap before the other is full. This priority queue also achieves the desired worst-case time and comparison bounds, but it uses n additional bits, so the structure is still not quite in-place.

For a long time, it was open whether there exists an in-place data structure that guarantees $O(1)$ worst-case time for *minimum* and *insert*, and $O(\lg n)$ worst-case time for *extract-min* such that *extract-min* performs at most $\lg n + O(1)$ element comparisons. In view of the lower bounds proved in [13], it was not entirely clear if such a structure exists. A strengthened lazy heap, proposed in this paper operates in-place and supports *minimum*, *insert*, and *extract-min* within the desired bounds, and thus settles this long-standing open problem.

When a strengthened lazy heap is used in heapsort, the resulting algorithm sorts n elements in-place in $O(n \lg n)$ worst-case time performing at most $n \lg n + O(n)$ element comparisons. The number of element comparisons performed matches the information-theoretic lower bound for sorting up to the additive linear term. Ultimate heapsort [14] is known to have the same complexity bounds, but the constant factor of the additive linear term is high.

In a binary heap the number of element moves performed by *extract-min* is at most $\lg n + O(1)$. We have to avow that, in our data structure, *extract-min* may require more element moves. On the positive side, we can adjust the number of element moves to be at most $(1 + \varepsilon) \lg n$, for any fixed constant $\varepsilon > 0$ and large enough n, while still achieving the desired bounds for the other operations.

One motivation for writing this paper was to show the limitation of the lower bounds proved by Gonnet and Munro [13] (see also [3]) in their important paper on binary heaps. They showed that $\lceil \lg \lg(n+2) \rceil - 2$ element comparisons are necessary to insert an element into a binary heap. In addition, slightly correcting [13], Carlsson [3] showed that $\lceil \lg n \rceil + \delta(n)$ element comparisons are necessary and sufficient to remove the minimum element from a binary heap that has $n > 2^{h_{\delta(n)}+2}$ elements, where $h_1 = 1$ and $h_i = h_{i-1} + 2^{h_{i-1}+i-1}$. One should notice that these lower bounds are valid under the following assumptions:

1) All elements are stored in one nearly complete binary tree.

2) Every node obeys the heap property before and after each operation.

3) No order relation among the elements of the same level can be deduced from the element comparisons performed by previous operations.

We show that the number of element comparisons performed by *extract-min* can be lowered to at most $\lg n + O(1)$ if we overrule the third assumption by imposing an additional requirement that any right child is not smaller than its left sibling; see Section 2. We also show that *insert* can be performed in $O(1)$ worst-case time if we overrule the second assumption by allowing $O(\lg^2 n)$ nodes to violate heap order; see Section 3. Lastly, we combine the two ideas and put everything together in our final data structure; see Section 4.

2 Strong Heaps: Adding More Order

A *strong heap* is a binary heap with one additional invariant: The element at any right child is not smaller than that at its left sibling. This left-dominance property is fulfilled for every right child in a fine heap [4] (see also [16, 18]), which uses one extra bit per node to maintain the property. By analogy, the heap property can be seen as a dominance relation between a node and its parent. As a binary heap, a strong heap is viewed as a nearly complete binary tree where the lowest level may be missing some nodes at the rightmost (last) positions. Also, this tree is embedded in an array in the same way. If the array indexing starts at 0, the parent of a node at index i ($i \neq 0$) is at index $\lfloor (i-1)/2 \rfloor$, the left child (if any) at index $2i+1$, and the right child (if any) at index $2i+2$.

Two alternative representations of a strong heap are exemplified in Fig. 1. On the left, the *directed acyclic graph* has a complete binary tree as its skeleton: There are arcs from every parent to its children and additional arcs from every left child to its sibling indicating the dominance relations. On the right, in the *stretched tree* the arcs from each parent to its right child are removed reflecting that these dominance relations can be induced. In the stretched tree a node can have 0, 1, or 2 children. A node has one child if in the skeleton it is a right child that is not a leaf or it is a leaf that has a right sibling. A node has two children if in the skeleton it is a left child that is not a leaf. If the underlying nearly complete binary tree has height h, the height of the stretched tree is at most $2h-1$ and

on any path from the root to a leaf in the stretched tree the number of nodes with two children is at most $h - 2$.

The basic primitive used in the manipulation of a binary heap is the *sift-down* [12, 19]. The procedure starts at a node that breaks the heap property, traverses down the heap by following the smaller children, and moves the encountered elements one level up until the correct place of the element we started with is found. In principle, *sift-down* performs a single iteration of insertion sort to get the elements on this path in sorted order. In a strong heap, the *strengthening-sift-down* procedure does the same, except that it operates on the stretched tree. Having this tool available, *extract-min* can be implemented by replacing the element at the root with the element at the last leaf of the array, and invoking the *strengthening-sift-down* procedure for the root.

Example 1 *Consider the strong heap in Fig. 1. If its minimum was removed, the special path to be followed would include the nodes* $\langle 3, 4, 5, 7, 11 \rangle$.

To build a strong heap for a given set of elements, we can imitate Floyd's heap-construction algorithm [12]; that is, *strengthening-sift-down* is invoked for all nodes, one by one, processing them in reverse order.

Let n denote the size of the strong heap under consideration, and let h denote the height of the underlying nearly complete binary tree. When we go down the stretched tree, we have to perform at most $h-2$ element comparisons (branching at binary nodes), and when we check whether to stop or not, we have to perform at most $2h - 1$ element comparisons. Hence, the number of element comparisons performed by *extract-min* is bounded by $3h - 3$, which is at most $3 \lg n$ knowing that $h = \lfloor \lg n \rfloor + 1$. The total number of element comparisons performed by *construct* is bounded by the sum $\sum_{i=1}^{\lfloor \lg n \rfloor + 1} 3 \cdot i \cdot \lceil n/2^{i+1} \rceil$, which is at most $3n + o(n)$. For both procedures, the amount of work done is proportional to the number of element comparisons performed, i.e. the worst-case running time of *extract-min* is $O(\lg n)$ and that of *construct* is $O(n)$.

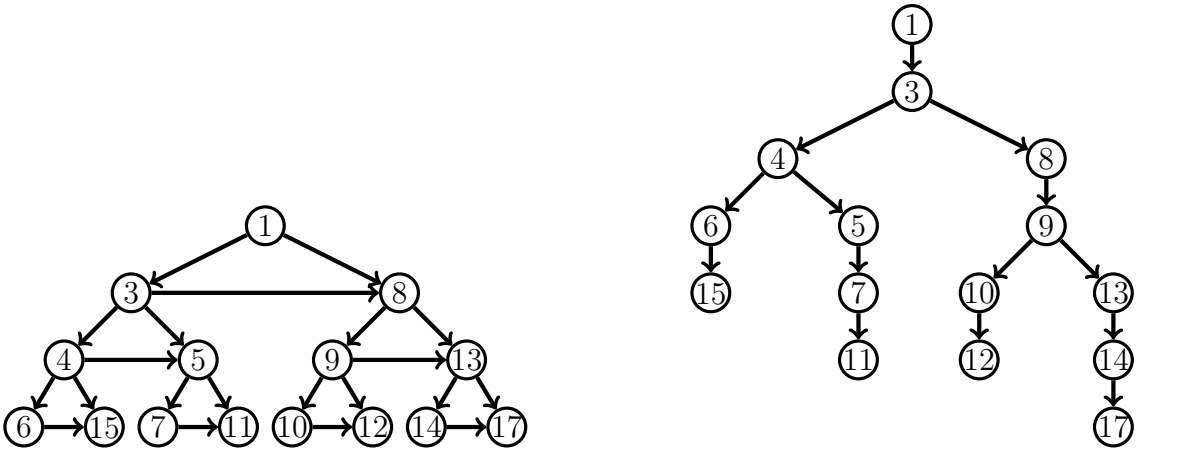

Figure 1: A strong heap in an array $a = [1, 3, 8, 4, 5, 9, 13, 6, 15, 7, 11, 10, 12, 14, 17]$ and its alternative representations: directed acyclic graph (left) and stretched tree (right)

Lemma 1 *A strong heap of size n can be built in $O(n)$ worst-case time by repeatedly calling strengthening-sift-down. One strengthening-sift-down operation uses $O(\lg n)$ worst-case time and performs at most $3 \lg n$ element comparisons.*

The next question is how to perform a *sift-down* operation on a strong heap of size n using at most $\lg n + O(1)$ element comparisons. At this stage we allow the amount of work to be much higher, namely $O(n)$. We show how to achieve these bounds when the heap is *complete*, i.e. when all leaves have the same depth.

To keep the heap complete, assume that the element at the root of a strong heap is to be replaced by a new element. The *swapping-sift-down* procedure traverses the left spine of the root of the skeleton bottom up starting from the leftmost leaf, and determines the correct place of the new element using one element comparison at each node visited. Thereafter, it moves all the elements above this position on the left spine one level up, and inserts the new element into this place. If this place is at level h, we have performed h element comparisons. Up along the left spine of the root there are $\lg n - h + O(1)$ remaining levels to which we have moved other elements. While this results in a heap, we still have to reinforce the left-dominance property at these upper levels. In accordance, we compare each element that has moved up with its right sibling. If the element at position j is larger than the element at position $j+1$, we have to interchange the subtrees T_j and T_{j+1} rooted at positions j and $j+1$ in-place by swapping all their elements. The procedure continues this way until the root is processed.

Example 2 *Consider the strong heap in Fig. 1. If the element at the root was replaced with element 16, the left spine to be followed would include the nodes $\langle$(3), (4), (6)$\rangle$, the new element would be placed at the leaf we ended up, the elements above would be lifted up one level, and the subtree interchange would be necessary for the subtrees rooted at node (6) and its new sibling (5).*

Given two subtrees that are both complete and of height h, the number of element moves needed to interchange the two subtrees is $O(2^h)$. As the sum of $O(2^h)$ over all $h = 1, \ldots, \lfloor \lg n \rfloor$ is $O(n)$, the total work done in the subtree interchanges is $O(n)$. Thus, the *swapping-sift-down* operation requires at most $\lg n + O(1)$ element comparisons and $O(n)$ work.

Lemma 2 *In a complete strong heap of size n, a single swapping-sift-down operation can be executed in-place using at most $\lg n + O(1)$ element comparisons. The number of element moves performed is $O(n)$.*

3 Lazy Heaps: Buffering Insertions

In the variant of a binary heap that we describe in this section some nodes may violate heap order, because insertions are buffered and unordered bulks are incrementally melded with the heap. The main difference between the present construction and the construction in

[8] is that, for a heap of size n, we allow $O(\lg^2 n)$ heap-order violations, instead of $O(\lg n)$, but still we use $O(1)$ extra space to track where the potential violations are. By replacing *sift-down* with *strengthening-sift-down* the construction will also work for strong heaps.

A *lazy heap* is composed of three parts: 1) *main heap*, 2) *submersion area*, and 3) *insertion buffer*. The following rules are imposed:

1) New insertions are appended to the insertion buffer at the end of the array.

2) The submersion area is incrementally melded with the main heap by performing a constant amount of work in connection with every *insert* operation.

3) When the insertion buffer becomes full, the submersion process must have been completed. A proportion of the elements of the insertion buffer are then treated as a new submersion area.

4) When the insertion buffer is empty, the submersion process must have been completed. When an *extract-min* operation is performed, a replacement element is thus borrowed from either the insertion buffer or the main heap (if the insertion buffer is empty) but never from the submersion area.

The insertion buffer occupies the last locations of the array. If the size of the main heap is n_0, the insertion buffer is never allowed to become larger than $\Theta(\lg^2 n_0)$. The buffer should support insertions in constant time, minimum extractions in $O(\lg n)$ time using at most $\lg n + O(1)$ element comparisons, and it should be possible to meld a buffer with the main heap efficiently.

Let $t = \lfloor \lg(1 + \lg(1 + n_0)) \rfloor$. We treat the insertion buffer as a sequence of *chunks*, each of size $k = 2^{t-1}$, and we limit the number of chunks to at most k. All the chunks, except possibly the last one, will contain exactly k elements. The minimum of each chunk is kept at the first location of the chunk, and the index of the minimum of the buffer is maintained. When this minimum is removed, the last element is moved into its place, the new minimum of that chunk is found in $O(k)$ time using $k - 1$ element comparisons (by scanning the elements of the chunk), and then the new overall minimum of the buffer is found in $O(k)$ time using $k - 1$ element comparisons (by scanning the minima of the chunks).

When an *extract-min* operation needs a replacement for the old minimum, we have to consider the case where the last element is the overall minimum of the insertion buffer. In such a case, to avoid losing track of this minimum, we start by swapping the last element of the insertion buffer with its first element.

In *insert*, a new element is appended to the buffer. Subsequently, the minimum of the last chunk and the minimum of the buffer are adjusted if necessary; this requires at most two element comparisons. Once there are k full chunks, the first half of them are used to form a new submersion area and the elements therein are incrementally inserted into the main heap as a bulk.

A *bulk-insert* procedure is performed incrementally by visiting the ancestors of the nodes containing the initial bulk as in Floyd's heap-construction algorithm [12]. During such a submersion, the submersion area is treated as part of the main heap whose nodes may not obey the heap property. To perform a bulk insertion, heap order is reestablished bottom up level by level. Starting with the parents of the nodes containing the initial bulk, for each node we call the *sift-down* subroutine. We then consider the parents of these nodes at the next upper level, restoring heap order up to this level. This process is repeated all the way up to the root. As long as there are more than two nodes that are considered at a level, the number of such nodes almost halves at the next level.

In the following analysis we separately consider two phases of the *bulk-insert* procedure. The first phase comprises the *sift-down* calls for the nodes at the levels with more than two involved nodes. Let b denote the size of the initial bulk. The number of the nodes visited at the jth last level is at most $\lfloor (b-2)/2^{j-1} \rfloor + 2$. For a node at the jth last level, a call to the *sift-down* subroutine requires $O(j)$ work. In the first phase the amount of work involved is $O(\sum_{j=2}^{\lceil \lg r \rceil} j/2^{j-1} \cdot b) = O(b)$. The second phase comprises at most $2\lfloor \lg n_0 \rfloor$ calls to the *sift-down* subroutine; this accounts for a total of $O(\lg^2 n_0)$ work. Since $b = \Theta(\lg^2 n_0)$, the overall work done is $O(\lg^2 n_0)$, i.e. amortized constant per *insert*. Instead of doing this job in one shot, we distribute the work by performing $O(1)$ work in connection with every heap operation. Obviously, the bulk insertion should be done fast enough so that it completes before the insertion buffer becomes either full or empty.

To track the progress of the submersion process, we maintain two intervals that represent the nodes up to which the *sift-down* subroutine has been called. Each such interval is represented by two indices indicating its left and right endpoints, call them (ℓ_1, r_1) and (ℓ_2, r_2). These two intervals are at two consecutive levels, and the parent of the right endpoint of the first interval has an index that is one less than the left endpoint of the second interval, i.e. $\ell_2 - 1 = \lfloor (r_1 - 1)/2 \rfloor$. We call these two intervals the *frontier*. Notice that the submersion area consists of all the descendants of the nodes at the frontier. While the submersion process advances, the frontier moves upwards and shrinks until it has one or two nodes. This frontier imparts that a *sift-down* is being performed starting from the node whose index is $\ell_2 - 1$. In addition to the frontier, we also maintain the index of the node that the *sift-down* in progress is currently processing. In connection with every heap operation, the current *sift-down* progresses a constant number of levels downwards and this index is updated. Once the *sift-down* operation returns, the frontier is updated. When the frontier reaches the root, the submersion process is complete. To summarize, the information maintained to record the state of a bulk insertion is two intervals of indices to represent the frontier plus the node which is under consideration by the current *sift-down* operation.

As for the insertion buffer, we also maintain the index of the minimum among the elements on the frontier. We treat each of the two intervals of the frontier as a set of consecutive chunks. Except for the first or last chunk on each interval that may have less nodes, every other chunk has k nodes. In addition, we maintain the invariant that the minimum within every chunk on the frontier is kept at the entry storing the first node

among the nodes of the chunk. An exception is the first and last chunks, where we maintain the index for the minimum on each.

To remove the minimum of the submersion buffer, we know it must be on the frontier and readily have its index. This minimum is swapped with the last element of the array and a *sift-down* is performed to remedy the order between the replacement element and the elements in its descendants. We distinguish between two cases: 1) If there are at most two nodes on the frontier, we make the minimum index of the frontier point to the smaller. 2) If there are more than two nodes on the frontier, the height of the nodes on the frontier is at most $2 \lg \lg n + O(1)$. The nodes of the chunk that contained the removed minimum are scanned to find its new minimum. If this chunk is neither the first nor the last of the frontier, its new minimum is swapped with the element at its first position, and this is followed by a *sift-down* performed on the latter element. The overall minimum of the frontier is then localized by scanning the minima of all the chunks. Extracting the minimum of the submersion buffer thus requires $O(\lg n)$ time using at most $\lg n + O(1)$ element comparisons.

4 Strengthened Lazy Heaps: Putting Things Together

Our final construction is similar to the one of the previous section in that it consists of three components: main heap, submersion area, and insertion buffer. The main heap has two layers: a *top heap* that is a binary heap, and each leaf of the top heap roots a *bottom heap* that is a complete strong heap. Because the main heap is only partially strong, we call the resulting data structure a *strengthened lazy heap*.

Let n_0 be the size of the main heap, and let $t = \lfloor \lg(1 + \lg(1 + n_0)) \rfloor$. The height of the bottom heaps is t, or possibly $t + 1$. In the insertion buffer, the size of a chuck is $k = 2^{t-1}$ and the size of the buffer is bounded by k^2.

Because of the two-layer structure, the incremental remedy processes are more complicated for a strengthened lazy heap than for a lazy heap. Let us consider the introduced complications one at a time. To help the reader to get a complete picture of the data structure, we visualize it in Fig. 2.

Complication 1. Due to insertions and extractions, we should be able to move the *border* between the two layers dynamically. To make the bottom heaps one level shallower, we can just adjust t and ignore the left domination for the children of the nodes at the previous border. To make the bottom heaps one level higher, we need a new incremental remedy process that scans the nodes on the old border and compares every left child with its right sibling. If the right sibling is smaller, the two elements are swapped and a *strengthening-sift-down* operation is incrementally applied on the right sibling. Again, we only need a constant amount of space to record the current state of this process. When moving the border upwards, the total work done is linear so, after the process is initiated, every forthcoming operation has to take a constant share of this work. The *extract-min* operations do not need to be aware of this process, as they can always make a conservative assumption about the level at which the border is.

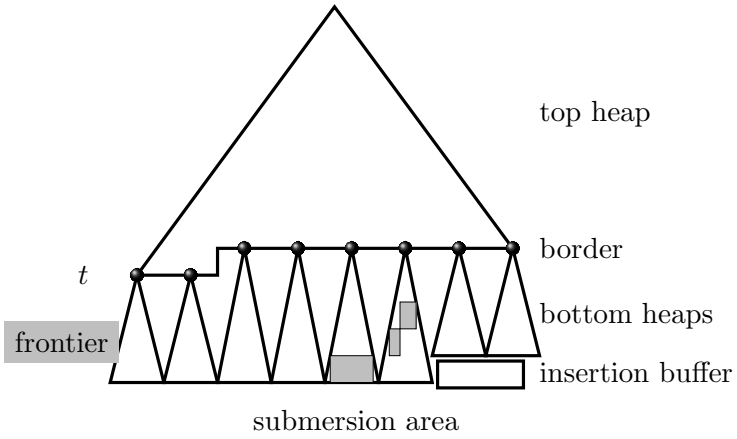

Figure 2: Schematic view of a strengthened lazy heap

Complication 2. We need a new procedure, that we call *combined-sift-down*, instead of *sift-down*. Assume we have to replace the minimum element that is in the top heap with another element. To reestablish the heap properties, we follow the proposal of Carlsson [2]: We traverse down along the path of the smaller children until we reach a root of a bottom heap. By comparing the replacement element with the element at the root of this bottom heap, we check whether the replacement element should land in the top heap or in the bottom heap. In the first case, we find the position of the replacement element using binary search on the traversed path; the path is traversed sequentially, but element comparisons are made in a binary-search manner. In the second case, we apply the *swapping-sift-down* operation for the root of the bottom heap.

Complication 3. Normally, we use the last element in the array as a replacement for the old minimum. However, if the insertion buffer is empty, meaning that the submersion process must have been completed, we need to use an element from the main heap as a replacement. In such a case, to keep the bottom heaps complete, our solution is to move all the elements at the lowest level of the last bottom heap (these are the last elements of the array) back to the empty insertion buffer. We consciously adjusted the parameters such that the number of elements in one chunk of the insertion buffer matches the number of elements at the lowest level of the last bottom heap. After such a move, the minimum of this chunk is not known. Fortunately, we would never need this minimum within the next k *extract-min* operations, as there are at least a logarithmic number of elements in the main heap that are smaller. In such a case, the minimum of this chunk is incrementally found within the upcoming k heap operations.

Complication 4. In the top heap the frontier causes no problem and is treated as before. However, if we perform a *swapping-sift-down* in a bottom heap whose frontier consists of two intervals, there is a risk that we mess up the frontier. In accordance, we schedule the submersion process differently: We process the bottom heaps under submersion one by one, and lock the bottom heap that is under consideration while performing *extract-min* operations. When the frontier overlaps the bottom heaps, it is cut into several pieces:

1) the interval corresponding to the unprocessed leaves of the initial bulk, 2) the two intervals (ℓ_1, r_1) and (ℓ_2, r_2) in the bottom heap under consideration, and 3) the interval of the roots of the bottom heaps that have been handled by the submersion process. Above the unprocessed interval and below the processed interval, *swapping-sift-down* operations can be performed without complications. Locking solves the problem for the bottom heap under consideration. However, in that bottom heap there are some nodes between the root and the frontier that are in no man's land since they are locked, but not yet included in the submersion process. These nodes will accordingly not be in order with the elements above or below. This is not a problem however, as none of these elements can be the minimum of the heap except after a logarithmic number of operations. Within such time, these nodes should have already been handled by the submersion process.

Complication 5. There should not be any interference between the remedy processes. One issue arises when a *swapping-sift-down* operation meets a node that an incremental *strengthening-sift-down* operation is processing. To avoid wrong decisions, it only requires one more check in the *swapping-sift-down* procedure if it meets such a node on the left spine of a bottom heap.

Let us now recap how the operations are executed and analyse their performance. Here we ignore the extra work done due to the incremental processes. Clearly, *minimum* can be carried out in $O(1)$ worst-case time by reporting the minimum of three elements: 1) the element at the root of the top heap, 2) the minimum of the insertion buffer, and 3) the minimum of the submersion area. As before, *insert* appends the given element to the insertion buffer and updates the minimum of the buffer if necessary. To perform an *extract-min* operation, we need to consider the different minima and remove the smallest among them.

Case 1. If the minimum is at the root of the top heap, we find a replacement for the old minimum and apply the *combined-sift-down* for the root. If we meet the frontier, we stop the operation before crossing it. Let n denote the total number of elements. The top heap is of size $O(n/\lg n)$ and the bottom heaps are of size $O(\lg n)$. To reach the root of a bottom heap, we perform $\lg n - \lg\lg n + O(1)$ element comparisons. If we have to go upwards, we perform $\lg\lg n + O(1)$ additional element comparisons in the binary search. On the other hand, if we have to go downwards, the *swapping-sift-down* operation needs to perform at most $\lg\lg n + O(1)$ element comparisons. In both cases, the work done is $O(\lg n)$.

Case 2. If the minimum of the insertion buffer is the overall minimum, it is removed as explained in the previous section. The operation involves $2k + O(1)$ element comparisons and the amount of work done is proportional to that number. Since we have set $k = 2^{t-1} = \frac{1}{2}\lg n + O(1)$, the minimum extraction here also requires at most $\lg n + O(1)$ element comparisons and $O(\lg n)$ work.

Case 3. If the frontier contains the overall minimum, we apply a similar treatment to that explained in the previous section with a basic exception. If there are more than two nodes on the frontier, the height of the nodes on the frontier is at most $2\lg\lg n + O(1)$. In this case, we use the *strengthening-sift-down* procedure in place of the *sift-down* procedure. If there are at most two nodes on the frontier, The frontier lies in the top heap. In this case,

we apply the *combined-sift-down* procedure instead. Either way, the minimum extraction requires at most $\lg n + O(1)$ element comparisons and $O(\lg n)$ work as well.

When executing the subtree interchanges, the number of element moves performed—even though asymptotically logarithmic—would still be larger than the number of element comparisons. Alternatively, one can control the number of element moves performed by the *swapping-sift-down* operation by adjusting the heights of the bottom heaps. If the maximum height of a bottom heap is set to $t - \lg(3/\varepsilon)$, for some small constant $0 < \varepsilon \leq 3$, the number of element moves performed by the *swapping-sift-down* operation will be bounded by $\varepsilon \lg n + O(1)$, while the aforementioned bounds for the other operations still hold.

5 Conclusions

We described a priority queue that 1) operates in-place, 2) supports *minimum* and *insert* in $O(1)$ worst-case time, 3) supports *extract-min* in $O(\lg n)$ worst-case time, and 4) performs at most $\lg n + O(1)$ element comparisons per *extract-min*. The data structure is asymptotically optimal with respect to time, and optimal up-to-additive-constant with respect to space and element comparisons.

The developed data structure is a consequence of many years' work. Previous breakthroughs can be summarized as follows: 1) break the $2 \lg n + O(1)$ barrier for the number of element comparisons performed per *extract-min* when *insert* takes $O(1)$ worst-case time [9], 2) achieve the desired bounds using $O(n)$ words of extra space [10], 3) reduce the amount of extra space to $O(n)$ bits [8], 4) achieve the desired bounds in the amortized sense [11]. And, in this paper, we developed an in-place structure that achieves the desired bounds in the worst case.

It is interesting that we could surpass the two lower bounds known for binary heaps [13] by slightly loosening the assumptions that are intrinsic to these lower bounds. To achieve our goals, we simultaneously imposed more order on some nodes by forbidding some left children to be larger than their right siblings and less order on others by allowing some nodes to possibly be smaller than their parents. Our last one-line tip is that it may not be necessary to maintain the heap property—exactly as defined for binary heaps—after every operation!

References

[1] Brown, M.R.: Implementation and analysis of binomial queue algorithms. SIAM J. Comput. 7(3), 298–319 (1978)

[2] Carlsson, S.: A variant of Heapsort with almost optimal number of comparisons. Inform. Process. Lett. 24(4), 247–250 (1987)

[3] Carlsson, S.: An optimal algorithm for deleting the root of a heap. Inform. Process. Lett. 37(2), 117–120 (1991)

[4] Carlsson, S., Chen, J., Mattsson, C.: Heaps with bits. Theoret. Comput. Sci. 164(1–2), 1–12 (1996)

[5] Carlsson, S., Munro, J.I., Poblete, P.V.: An implicit binomial queue with constant insertion time. In: Karlsson, R., Lingas, A. (eds.) SWAT 1988. LNCS, vol. 318, pp. 1–13. Springer, Heidelberg (1988)

[6] Chen, J., Edelkamp, S., Elmasry, A., Katajainen, J.: In-place heap construction with optimized comparisons, moves, and cache misses. In: Rovan, B., Sassone, V., Widmayer, P. (eds.) MFCS 2012. LNCS, vol. 7464, pp. 259–270. Springer, Heidelberg (2012)

[7] Dutton, R.D.: Weak-heap sort. BIT 33(3), 372–381 (1993)

[8] Edelkamp, S., Elmasry, A., Katajainen, J.: Weak heaps engineered. J. Discrete Algorithms 23, 83–97 (2013), Presented at IWOCA 2012

[9] Elmasry, A.: Layered heaps. In: Hagerup, T., Katajainen, J. (eds.) SWAT 2004. LNCS, vol. 3111, pp. 212–222. Springer, Heidelberg (2004)

[10] Elmasry, A., Jensen, C., Katajainen, J.: Multipartite priority queues. ACM Trans. Algorithms 5(1), Article 14 (2008)

[11] Elmasry, A., Katajainen, J.: Towards ultimate binary heaps. CPH STL Report 2013-1, Department of Computer Science, University of Copenhagen (2013)

[12] Floyd, R.W.: Algorithm 245: Treesort 3. Commun. ACM 7(12), 701 (1964)

[13] Gonnet, G.H., Munro, J.I.: Heaps on heaps. SIAM J. Comput. 15(4), 964–971 (1986), Presented at ICALP 1982

[14] Katajainen, J.: The ultimate heapsort. In: CATS 1998. Australian Computer Science Communications, vol. 20, pp. 87–95. Springer, Singapore (1998)

[15] Knuth, D.E.: Sorting and Searching, The Art of Computer Programming, vol. 3. Addison Wesley Longman, Reading, 2nd edn. (1998)

[16] McDiarmid, C.J.H., Reed, B.A.: Building heaps fast. J. Algorithms 10(3), 352–365 (1989)

[17] Vuillemin, J.: A data structure for manipulating priority queues. Commun. ACM 21(4), 309–315 (1978)

[18] Wegener, I.: The worst case complexity of McDiarmid and Reed's variant of Bottom-up Heapsort is less than $n \log n + 1.1n$. Inform. and Comput. 97(1), 86–96 (1992), Presented at STACS 1991

[19] Williams, J.W.J.: Algorithm 232: Heapsort. Commun. ACM 7(6), 347–348 (1964)